Review of methods, progress and cross-validation studies
pertaining to
population trend and risk assessment
for Columbia River salmonids

E. E. Holmes
Northwest Fisheries Science Center
National Marine Fisheries Service
2725 Montlake Blvd. East
Seattle, WA 98112-2097

November 2003

BACKGROUND FOR THIS WHITE PAPER

The 2000 Federal Columbia River Power System (FCRPS) Biological Opinion (FCRPS Biop) evaluated whether the operation of the FCRPS, when combined with survival rates expected to occur in all other life stages of ESA listed salmonids, would result in a "high likelihood of survival and a moderate-to-high likelihood of recovery." This qualitative determination was informed by quantitative estimates for several evolutionarily significant units (ESU). Specifically, NOAA Fisheries evaluated:

- whether or not there would be a 5% or lower probability of absolute extinction of natural spawners within 24- and 100-year periods as a "metric indicative of survival;"
- whether or not there would be at least a 50% probability of the 8-year geometric mean natural spawners being equal to, or greater than, interim recovery abundance levels in 48 and 100 years as a primary "metric indicative of recovery;"
- and whether or not there would be at least a 50% likelihood of the annual population growth rate ("lambda") being equal to, or greater than, 1.0 as an alternate "metric indicative of recovery" for populations lacking interim recovery abundance goals.

The basis for each of these indicator metrics was an analysis of the population growth rate associated with time series for relevant spawning aggregations. Population growth rate was calculated using the methods described in McClure et al. (2003). The Biological Opinion specified that several tests based on population growth rate would be conducted in 2005 and 2008 to ensure that implementation of the Biological Opinion was on track and that populations were not declining further. The Biological Opinion assumed that by 2005 there would be more information about methods of calculating population growth rate, so it specified the following:

"NMFS anticipates that methods of assessing annual population growth rates will have been refined, based on NMFS' research efforts, those of the Action Agencies, or those of independent scientists. In anticipation of this normal progress in scientific methods, NMFS does not now define a specific method by which population growth rate will be determined for its mid-point evaluations. By March 1, 2005, NMFS will choose the most appropriate method(s) to estimate population growth rate from the peer-reviewed literature, based on collaboration with the Action Agencies, USFWS, and the state and Tribal comanagers."

In June 2003, the Biological Opinion was remanded in National Wildlife Federation v. NMFS. NOAA Fisheries is currently in the process of revising the Biological Opinion and re-evaluating the effects of FCRPS operations and offsite mitigation activities. To facilitate this process, the NOAA Fisheries Northwest Regional Office (NWRO) requested that the Northwest Fisheries Science Center (NWFSC) conduct the above-referenced review of population growth rate estimation methods in 2003. In addition, the NWRO requested that that the NWFSC review related methods of characterizing population trends, especially those that had been suggested as alternatives to "lambda" estimation in comments on the draft of the original Biological Opinion and in comments or litigation since the Biological Opinion was issued.

E. E. Holmes, National Marine Fisheries Service

INTRODUCTION

The purpose of this report is to review and discuss methods for estimating and presenting population trends and extinction risks for Columbia River salmonid populations to support management decisions, such as the ESA Section 7 determination in the 2000 FCRPS Biological Opinion and the anticipated 2005 and 2008 check-in tests. This report reviews research since 2000, which tests and validates diffusion approximation methods for estimating population trends and risks. This review summarizes information from the following publications:

Holmes, E. E. 2004. Beyond theory to application and evaluation: diffusion approximations for population viability analysis. In press in Ecological Applications.
Fagan, W. F., E. E. Holmes, J. J. Rango, A. Folarin, J. A. Sorensen, J. E. Lippe, and N. E. McIntyre. 2003. Cross-validation of quasi-extinction risks from real time series: an examination of diffusion approximation methods. Pre-print.
McClure, M., E. Holmes, B. Sanderson, and C. Jordan. 2003. A large-scale, multi-species risk assessment: anadromous salmonids in the Columbia River Basin. Ecological Applications 13: 964-989.
Holmes, E. E. and W. F. Fagan. 2002. Validating population viability analysis for corrupted data sets. Ecology 83: 2379-2386
Holmes, E. E. 2001. Estimating risks in declining populations with poor data. Proceedings of the National Academy of Science 98: 5072-5077.

Summary of work and changes as they pertain to the methods in FCRPS Biop

Changes in the methods for estimating trend and risk:
1) Running sum filter has been standardized to use a simple sum of four consecutive spawner counts. The work leading up to Holmes & Fagan (2002) clarified that this was better than the age-structure based running sum that was originally used.

Cross-validation work
The bulk of the work has focused on validating the methods using real time series (Holmes & Fagan 2002, Fagan et al. 2003) and more realistic simulations that include density-dependence (Holmes 2004). Also the underlying assumptions of the diffusion model were tested using simulations of salmon models with density-dependence (Holmes 2004).

Expressing uncertainty
Holmes & Fagan (2002) test the variability in parameter estimates from the Dennis-Holmes method and found that the variability is properly estimated. Holmes (2004) begins looking in-depth at how to express uncertainty in a way that it can best inform regulatory decision-making. Confidence intervals are commonly given, but are not very useful beyond showing that there is high or low uncertainty. Bayesian approaches are explored in Holmes (2004). A Bayesian metric is also used in McClure et al. (2003), specifically the probability that λ is less than 1.0 or less than 0.9.

E. E. Holmes, National Marine Fisheries Service

THE NATURE OF POPULATION TRAJECTORIES AND RISK ESTIMATION

Real populations do not grow or decline at fixed rates, but rather show year-to-year variability in population growth rates, which leads to a population trend that varies about some long-term growth rate. Figure 1 shows an example of three population trends that each have the same long-term trend (5% per year decline) and the same year-to-year variability. Even though the population trends were generated with the same underlying dynamics, the trajectories are different. This is nature of populations: random chance means that there are a range of different possible population trajectories given some underlying population dynamics. Even though we cannot predict exactly what will happen in the future, if we could estimate the underlying dynamics governing the population trajectories, we could estimate the probability of different futures, i.e. we could estimate the probability of reaching critical thresholds. We can also estimate whether the population has long-term declining dynamics. To do this, we will need to estimate the following: the long-term rate of decline (or growth), the year-to-year variability in yearly population growth, and the amount of corruption in our data. Within the population dynamics literature, the year-to-year variability in yearly population growth is termed 'process error'; note that it is not technically 'error' in the layman's sense of the word, but rather variability. The rest of the variability is termed 'non-process error' and this includes actual observations errors. For the purpose of this review, one can think of process error as the variability that drives the long-term variability of future population size and the non-process error as the data corruption that is preventing us from estimating the process error.
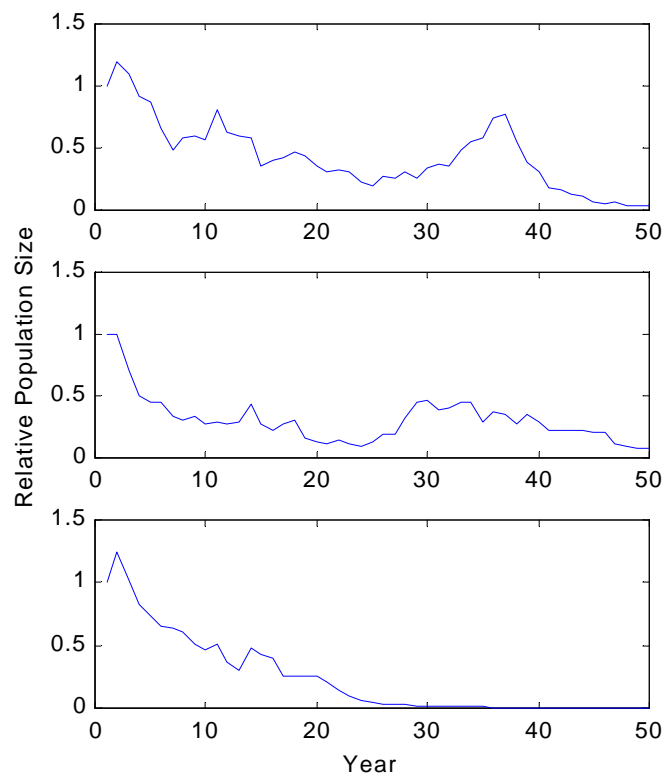

Figure 1. Sample simulated population trajectories from a population with average 5% yearly decline. The underlying population dynamics are identical. The differences are due to chance.

E. E. Holmes, National Marine Fisheries Service

The term λ denotes the long-term rate of population decline (or growth). It is simply the long-term trend that you would observe if you had a very, very long time series of the population. The term λ is the standard notation in the conservation biology literature. λ = 1 means stable, λ = 1.01 means roughly increasing 1% per year, and λ = 1.05 means roughly increasing 5% per year. Similarly, λ = 0.99 and λ = 0.95 mean roughly declining 1% and 5% per year, respectively. Note that we can only estimate λ; we never know the true λ. Our estimates may be unbiased, but that still means that there is a 50/50 chance that the true λ is above or below our estimate.

One of the most common questions is "If λ is the population trend, why not just present the overall trend observed in the data, such as a regression of log numbers versus time?" as opposed to going through the analysis based on theory concerning the dynamics of population trajectories, which is presented in the next section. There are two main reasons why this is insufficient.

1) *We need to estimate uncertainty*. The trend tells you what happened but does not by itself tell you how likely it is that this trend happened by chance and that the long-term dynamics are actually quite different. For example, suppose we collect data on a population that has a true long-term average rate of decline of 12% yearly. Figure 2 shows an example of the population trend observed from 20-year consecutive time series from this population. Segment 1 is from year 1-20, segment 2 is from year 2-21, etc. The wavy lines show the estimates using different methods for estimating the trend; the true value is the straight dashed line. The solid line ("ML") is a simple regression of log natural abundance. The wavy dashed line shows the runsum method used in McClure et al. (2003) and the Biop. There is much variability in the observed trend in a 20-year segment. This variability is an unavoidable aspect of analyzing stochastic population processes. Population dynamics theory allows us to estimate this year-to-year variability and thus estimate how likely it is that a particular observed trend came from a population with a particular true λ (such as an increasing or declining population). But to do this, the estimate of the underlying process error in the population dynamics is needed. A natural response would be to argue that standard regression analyses will give you the uncertainty of the estimated trend, but unfortunately such analyses attribute all error to non-process error and will give you incorrect uncertainty estimates.

2) *We need to estimate probabilities of crossing critical thresholds*. The trend by itself does not give much information about the probability of dropping below critical population sizes. We cannot simply extend the trend into the future and see when our line crosses the threshold. Populations vary from year to year and even a population that has a positive growth rate still has some probability of dropping below the critical threshold by chance. To estimate this probability, we again need to estimate the process error driving the variability in long-term population sizes.

In the following section (section I), I review how the parameters driving a population process are estimated using diffusion approximation methods. This section also reviews the extensive cross-validation work that was done to verify the applicability of these methods for salmon populations. This section directly applies to the methods used in the FCRPS Biological Opinion. At the end of this section, there is a discussion of alternative risk estimation methodologies and why they were not used. The next section (section II) discusses work that goes beyond the methods used in the 2000 FCRPS Biological Opinion. One of the challenges when presenting scientific analyses is presenting the uncertainty in a useful and accurate manner. It is tempting to use the point estimates of risk metrics (i.e. 'this stock has a λ of 0.981') and

E. E. Holmes, National Marine Fisheries Service

ignore that this is a statistical estimate. $\lambda = 0.981$ may be the most likely value given the data, but $\lambda = 0.99$ is probably almost equally as likely and $\lambda = 1.01$ may be entirely plausible. Section II illustrates the use of probability curves as a way to formally express this uncertainty. This is a standard approach in decision theory for resource management.
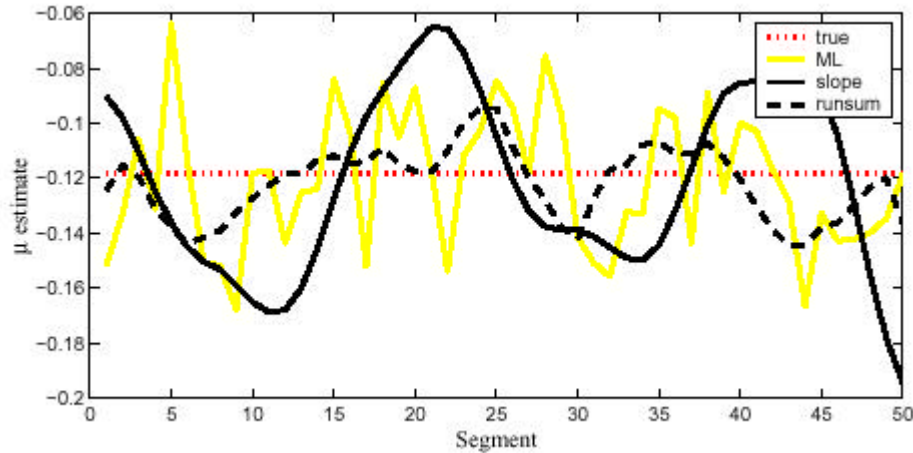

Figure 2. Estimated log(λ) from 20-year segments in a time series. Segment 1 is year 1-20, segment 2 is year 2-21, segment 3 is year 3-22, etc. This shows how the estimates vary depending on the segment observed. The "runsum" method is that used in McClure et al. (2003) and for the Biological Opinion calculations.

I. DIFFUSION APPROXIMATION METHODS FOR POPULATION VIABILITY ANALYSIS

In the last decade, diffusion approximation (DA) methods have been developed that use count data alone (for example, spawner counts) for the estimation of population viability analysis (PVA) risk metrics, such as the probability of crossing extinction thresholds, mean passage times, and average long-term rates of population growth or decline (Lande and Orzack 1988, Dennis et al. 1991). These methods have since been used to estimate extinction risks for numerous species of conservation concern (Dennis et al. 1991, Nicholls et al. 1996, Gerber et al. 1999, Morris et al. 1999, McClure et al. 2003). The appeal of DA methods from an applied standpoint is their simplicity and their reliance on simple census data alone (e.g. neither age-structure, cohort-level analyses, or total fish numbers are required). They have become one of the basic quantitative tools presented in recent books on PVA methods (Morris and Doak 2002, Lande et al. 2003).

Diffusion approximation methods stem from theory concerning the behavior of stochastic age-structured population models with no density-dependence,

$$\begin{bmatrix} n_{1,t+1} \\ n_{2,t+1} \\ n_{3,t+1} \\ \vdots \\ n_{k,t+1} \end{bmatrix} = \mathbf{A}_t \begin{bmatrix} n_{1,t} \\ n_{2,t} \\ n_{3,t} \\ \vdots \\ n_{k,t} \end{bmatrix} \qquad [1]$$

where $\mathbf{A}_t$ is the stochastic population transition matrix, e.g. a Leslie matrix, for time t. Note that most types of cohort or otherwise age-structured population simulations with no density-

dependence are specific cases of the general model in Eq. 1. For such models, the asymptotic behavior of the total population size, $N_t = \sum_i n_{i,t}$, is a stochastic exponential process (Tuljapurkar and Orzack 1980, Tuljapurkar 1989):

$$N_t = N_0 \exp(\boldsymbol{m}t + \boldsymbol{e}_p), \qquad \boldsymbol{e}_p \sim \text{normal}(0, \boldsymbol{s}^2 t) \text{ for } t \text{ big}. \qquad [2]$$

and $\log N_t/N_0$ is distributed normal with mean=μt and variance=$\sigma^2 t$ for t big. The parameter μ in Eqn. 2 determines the rate at which the median log population size, $\log N_t$, increases through time, while σ^2 determines the rate at which the distribution spreads, or in other words, the variability of potential population sizes at time $t+\boldsymbol{t}$.

  Diffusion approximation methods assume that Eqn. 2 holds for all $\tau > 0$ including small τ and that the ε are independently and identically distributed. This allows one to model the population as a diffusion process (Lande and Orzack 1988):

$$\partial p / \partial t = -\boldsymbol{m}\partial p / \partial x + (\boldsymbol{s}^2/2)\partial^2 p / \partial x^2$$
$$p = P(\log N_t = x \mid N_0 = x_0) \qquad [3]$$

P(y) means the probability of y. The diffusion model has the property that $\log N_t/N_0$ is distributed normal(mean=μt, variance=$\sigma^2 t$) like the stochastic exponential process it is used to approximate. See Dennis et al. (1991) for a much fuller discussion of the diffusion approximation.

  This approximation opens a toolbox of parameterization methods for linear models with normal error. It also provides analytical estimates of quasi-extinction probabilities, i.e. the probability of crossing a particular threshold at some time within a given time frame. Strictly speaking, however, an age-structured population process is not a diffusion process. However despite the assumption violations, the diffusion model approximates many types of stochastic age-structured population processes, as seen both from simulated and real data (Lande and Orzack 1988, Dennis et al. 1991, Holmes and Fagan 2002, Fagan et al. unpublished manuscript, Holmes 2004). In particular as will be reviewed below, the diffusion approximation works well for salmon population models (Holmes 2004).

Parameter estimation methods

  Diffusion approximations for a particular PVA must be carefully selected since a poor choice results in poor, highly biased estimates which lead to poor, highly biased risk estimates. Holmes (2004) discusses these issues and careful selection of parameterization methods using salmon data as an example. McClure et al. (2003) presents methods for estimating log(λ) and σ^2, which have been used by NWFSC scientists for salmon PVA. These methods have been extensively validated with real and simulated salmon data (see next section).

  The basic estimation methods currently used for the Biop are presented here without discussion; see McClure et al. (2003) for a discussion and examples. The methods use a running sum transformed time series of spawner counts defined as

$$R_t = \sum_{i=0}^{3} O_{t+i} \qquad [4]$$

where O_t is the spawner count at year t. The estimate for log(λ), which is denoted μ, is

$$\hat{\boldsymbol{m}}_{run} = \text{mean of } \log(R_{t+1}/R_t)$$
$$\text{for } t = 1,2,3,\ldots,k-3. \qquad [5]$$

The estimate of σ^2 uses the rate that the variance increases within the time series:

E. E. Holmes, National Marine Fisheries Service

$$\hat{s}_{slp}^2 = \text{slope of } \text{var}\left[\log\left(\frac{R_{t+t}}{R_t}\right)\right] \quad \text{versus } t, \text{ intercept free} \tag{6}$$

for $t = 1,2,3,\ldots, k-3$ and maximum $t = 4$.

These estimators will likely appear somewhat peculiar on first introduction. Note that the μ estimate is very similar to a linear regression of log population counts (typically log spawner counts). Why use the estimator with a running sum transformation of the data? Extensive testing described in Holmes (2001), Holmes & Fagan (2002) and especially Holmes (2004) indicates that the $\hat{m}_{run}$ gives the least variable estimates of μ (see also Figure 2). Estimation of the process error is an especially difficult problem. Holmes (2004) reviews the currently available methods in the literature. Again extensive cross-validation work (see especially Holmes 2004) found that $\hat{s}_{slp}^2$ performs the best for salmon data.

One of the difficult problems with analyzing salmon spawner data is that hatchery fish are input into the stocks. Perhaps the easiest way to see how this presents a problem for estimating λ is to consider the analogy of a mutual fund. Suppose you put $1000 into a mutual fund 5 years ago and now you have $8000. You would like to know what the average rate of return (this is λ) has been so that you can decide whether to keep you money in this fund or move to another. Normally, you would just take $(8000/1000)^{\wedge}(1/5) = 1.51$, which means that your fund returned an incredible 51% per year. However, your benevolent aunt has been automatically adding $100 a month to your brokerage account, and you need to factor this in (these are the hatchery fish). Problem is you don't know whether her monthly gift was added to your mutual fund (the hatchery fish reproduce) or was simply deposited to your brokerage account but not invested (the hatchery fish don't reproduce). Without this information, you can only deduce the range of the possible average rates of return for your mutual fund. If not added to mutual fund, the rate of return was $(\,(8000-100*12*5)/1000\,)^{\wedge}(1/5) = 1.15$, which is still a nice 15% per year. If added to the mutual fund, rate of return is found by finding the λ that solves:

$$8000 = 1000*\lambda^{\wedge}5 + 100*\sum_{i=0}^{12*5} l^{i/12},$$

which is $\lambda = 1.05$ and means a rather paltry 5% per year growth. Thus, knowing whether the monthly deposits were added to your mutual fund or not is a critical bit of information you need to evaluate how good a mutual fund you have. This is exactly the estimation problem we have with hatchery fish. We need to know whether or not they are reproducing in order to evaluate the underlying population growth rate. In McClure et al. (2003) and in Holmes (2004), the hatchery correction is presented. In the McClure et al. PVA, the range of λ for hatchery fish not reproducing versus are reproducing is shown. For the Biological Opinion, the range of λ is shown for hatchery fish reproducing 20% as effectively as wild-born fish versus 80% as well as wild-born fish.

Risk metrics

From the parameters μ and σ^2, a number of different risk metrics can be calculated. We have focused on two metrics. The first is the median yearly growth rate or the long-term yearly growth rate, which is denoted λ. Suppose you were able to observe 1000 20-year population trajectories with the same underlying dynamics (i.e. the same μ and σ^2 parameters) and each

starting from the same initial size, much like Figure 1. The trajectories would all look different due to chance. The yearly growth rate you observed in the i^{th} (out of the 1000) trajectory is

λ_i = [(end population size)/(start population size)]^(1/(number of years)-1)

The median λ_i from all 1000 would be $\exp(\mu)$; on average 50% of trajectories would show a yearly growth rate greater than $\exp(\mu)$ in those 20 years and 50% would show a lower growth rate. An estimate of this median yearly growth rate is what we term λ. It also happens to be the yearly growth you would observe from a very long time series since λ_i goes to $\exp(\mu)$ as the number of years gets very large. This is why the λ estimate is referred as an estimate of the median yearly growth rate or the long-term yearly growth rate. For a particular time series with n years, the λ estimate is

$$\hat{I} = \exp(\hat{m}_{run}) = \frac{1}{n-4}\log(R_{n-3}/R_1) \qquad [7]$$

The second metric is the probability of hitting a particular critical population threshold, N_e, within some period of time t_e, starting from the population size N_0. This is calculated from a diffusion approximation of the population process (Dennis et al. 1991):

$$\Pr(N_0 \rightarrow N_e \quad by \quad t_e) = \boldsymbol{p}' * \Phi\left(\frac{-\ln(N_0/N_e) + |\hat{\boldsymbol{m}}_{run}|\ t_e}{\hat{\boldsymbol{s}}_{slp}\sqrt{t_e}}\right)$$

$$+ \exp(2\ln(N_0/N_e)|\hat{\boldsymbol{m}}_{run}|/\hat{\boldsymbol{s}}_{slp}^2)\Phi\left(\frac{-\ln(N_0/N_e) - |\hat{\boldsymbol{m}}_{run}|\ t_e}{\hat{\boldsymbol{s}}_{slp}\sqrt{t_e}}\right) \quad t_e > 0 \qquad [8]$$

$$\text{where} \quad \boldsymbol{p}' = \begin{cases} 1, & if \quad \hat{\boldsymbol{m}}_{run} \leq 0 \\ \exp(-2\hat{\boldsymbol{m}}_{run}\ln(N_0/N_e)/\hat{\boldsymbol{s}}_{slp}^2), & if \quad \hat{\boldsymbol{m}}_{run} > 0 \end{cases}.$$

The function Φ is the standard normal cumulative distribution function. If you are interested in percentage-wise declines, e.g. 50%, 75% or 90%, then it is not necessary to know the actual population size, since $(N_0/xN_0) = (1/x)$. In this case, the probability of crossing critical thresholds can be estimated using on index data without information on the total number of spawners. If however, declines to specific absolute thresholds are of interest, total spawner counts are needed and it is also necessary to transform the spawner count into a count that reflects the total population rather than just spawners in a particular year. See McClure et al. (2003) for a discussion of this transformation.

Validation studies of diffusion approximations for salmon populations

Here I review cross-validation studies of the performance of the diffusion approximation model for salmon data and populations, including populations experiencing density dependence. Holmes (2004) discusses evaluation of the diffusion approximation and estimation methods using simulated data. This study used detailed population models for Upper Columbia River steelhead, Snake River fall chinook, and Snake River spring/summer chinook as examples. The models were parameterized from survivorship and fecundity data from these ESUs. The models include density-dependence in parr to smolt survivorship reflecting that found in low density Snake River chinook stocks (Achord et al. 2003). The models also include sampling error in the range of that observed for redd-count data (standard error 0.3 to 0.85).

The first question in this study was whether a diffusion approximation correctly described the behavior and probability of crossing thresholds for the age-structured models. The first test

described in Holmes (2004) is an examination of the linearity assumptions inherent in the diffusion approximation. This key test is somewhat technical, and is described in Holmes (2004). The results of this test were that the linearity assumptions were satisfied for $t > 5$ years which means that (as is well-known) the diffusion approximation should be used to make medium and long-term projections not short-term projections ($t < 5$ years). The second test was whether a diffusion approximation would properly characterize the probability that the simulated time series would cross a threshold (in this case, 90% decline) in different time frames. This analysis is shown in Figure 3. The gray line shows the actual probability of crossing the 90% decline threshold within different time frames (determined by repeating the salmon simulations 1000s of times) versus the probabilities from a diffusion approximation. This illustrates that the probability of 90% decline in these salmon time series can be described by a diffusion approximation.
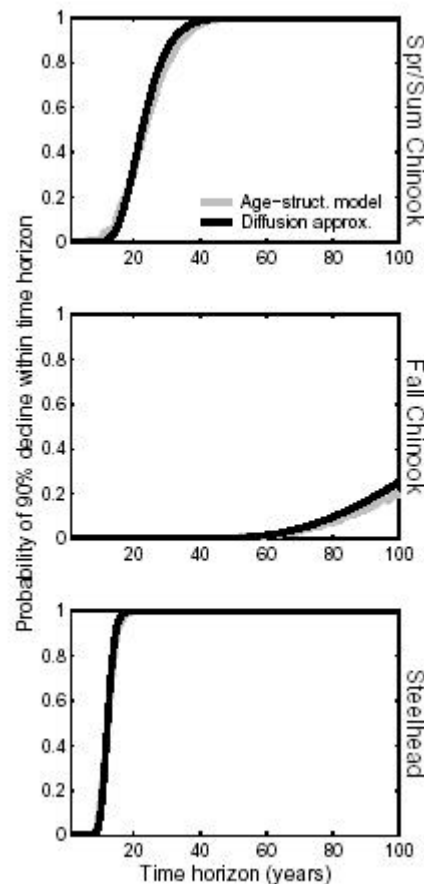

Figure 3. Actual versus predicted probability of 90% decline within different time horizons. From Holmes (2004).

Simply because a diffusion approximation exists which properly characterizes a particular population process does not mean that we can estimate the parameters for that process given realistic data constraints. Holmes (2004) also studies estimation performance given data constraints faced by the PVA of Columbia River salmon stocks (McClure et al. 2003): 1) counts of only the spawning segment of the populations, 2) time series limited to 20 years, 3) severe age-structure perturbations in the beginning of some time series due to reproductive collapses during dam construction (Williams et al. 2001), and 4) high observation error in the spawner

counts. Figure 4 shows box plots of the estimates of log(λ) following the estimation methods described above (also in McClure et al. 2003) for 1000 random simulations from the three species' models. The output from the models (spawner counts) was 'corrupted' by different levels of sampling error: age (meaning an age perturbation due to no reproduction in one year), low, medium and high observer error. The dotted line in the graph shows the true value of log(λ). In the box plots, the middle line is the median estimate of log(λ) and the box encloses 75% of the estimates. As can be seen in the figure, the runsum method for estimating log(λ) works for these simulated salmon time series even within the data constraints; the median estimate is the true value even with added observer error in the spawner counts.
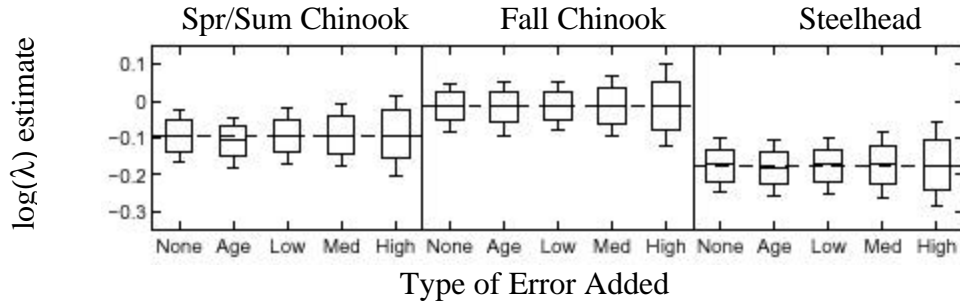

Figure 4. Distribution of log(λ) estimates using $\hat{\boldsymbol{m}}_{run}$ from 1000 simulated time series from age-structured models of Snake River spring/summer chinook, Snake River fall chinook, and Upper Columbia steelhead. The models include density-dependent smolt survivorship. From Holmes (2004).

Figure 5 shows a similar analysis for the estimation of the process error, termed σ^2. Recall that the process error specifies the variability of potential future population trajectories and is a key parameter determining the probability of crossing thresholds. This analysis indicates that for low observation error $\hat{\boldsymbol{s}}^2_{slp}$ provides an unbiased estimate of the true value of σ^2, but as observation error increased, the estimate becomes increasingly biased. 'Medium' represents the average estimate of typical observation error in the Columbia River data based on studies of observer error in redd count data (see discussion in Holmes 2004).
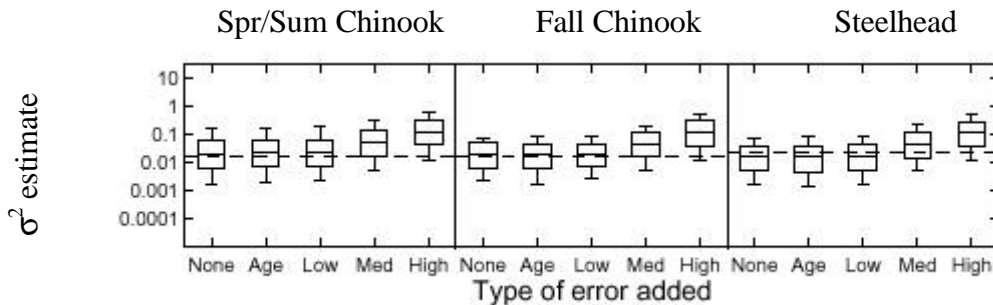

Figure 5. Distribution of $\hat{\boldsymbol{s}}^2_{slp}$ estimates from 1000 simulated time series from age-structured models of Snake River spring/summer chinook, Snake River fall chinook, and Upper Columbia steelhead. The models include density-dependent smolt survivorship. From Holmes (2004).

Simulated data is very useful, however it is 'simulated' and certainly lacks some aspects of real time series data. Another cross-validation (Holmes and Fagan 2002) involved testing the bias and precision of the diffusion approximation parameter estimates using hundreds of real time series. The strategy was to use the first 15 years of a time series to predict the second 15 years of the time series. The bias and variability of these predictions could then be tested against the predicted bias and variability. The two parameters tested were $\log(\lambda)$ and σ^2, which appears in the probability of crossing thresholds metric along with $\log(\lambda)$. Figure 6 shows the results of this analysis for the $\log(\lambda)$ estimates. This analysis involved 30-year time series within a 1920 to 1999 time frame.
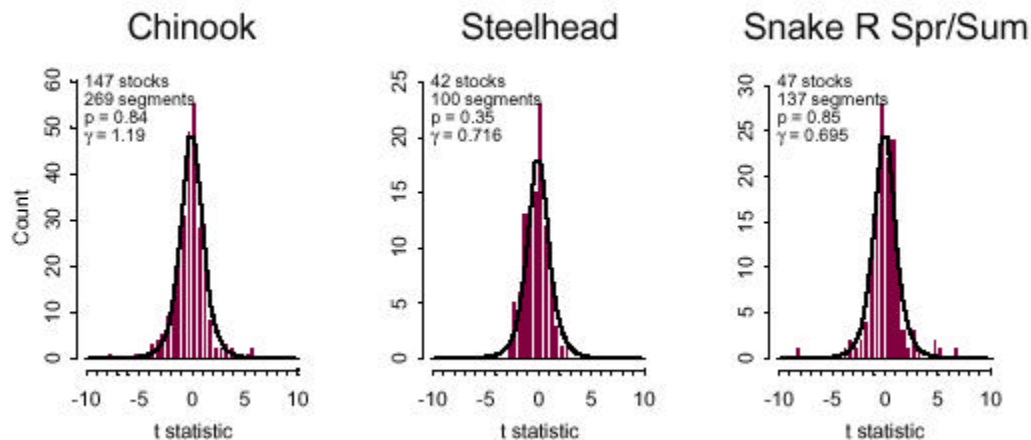

Figure 6. Distribution of actual $\log(\lambda)$ estimates (bars) versus predicted distribution (solid black line from 147, 42 and 47 chinook, steelhead and Snake R spring/summer time series respectively. From Holmes and Fagan (2002).

The close match between the observed and predicted distributions indicates that the $\log(\lambda)$ estimate was properly characterized in terms of its mean value (the peaks match). That is the mean trend in the first half of the time series was the same as the mean trend estimated in the second half of the time series. Figure 6 also demonstrates that the uncertainty in the $\log(\lambda)$ estimate (its variability) was also properly characterized since the width of the distributions match.

That the mean trend in the first 15 years was the same as the mean trend in the second 15 years appears at first glance to contradict the observations of strings of 'good years' versus 'bad mind that this analysis used 30-year time series across the 1920 to 1999 period. It was asking about the average estimate across different time periods. What about estimates only during a specific time period? Figure 7 shows the difference between the trend in the first 15 years of a time series versus the following 15 years for specific time frames, i.e., not the average across all time periods, but the average if you only look at time series in a specific time period, say 1970-1999. The solid line is a measure of the difference between the trends in the first 15 years versus the following 15 years. Deviations above zero indicate that on average there was a more declining trend in the first 15 years versus the next 15; while deviations below zero means that on average the population was declining less in the first 15 years versus the next 20 years. These results show the average difference from all the West Coast time series put together. What you can see is that across the West Coast, stocks were on average declining more in 1959-1973 versus in 1974-1993 while the opposite was true for 1964-1978 versus 1979-1998.

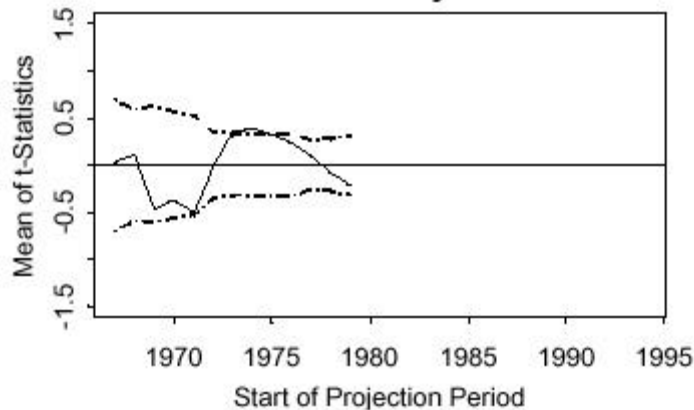E. E. Holmes, National Marine Fisheries Service

Figure 7. The solid line measures the difference between the average trend in the first 15 years versus in the second 15 years from a collection of 30 year time series of West Coast chinook and steelhead stocks (200+ stocks). "Mean of t-statistics" = 0 means that the average trend (across the whole West Coast) was similar in the first 15 years versus the following 15 years. "Mean of t-statistics" > 0 means that the on average stocks were declining more in the first 15 years relative to the following 15 years. The year on the x-axis denotes the start of the middle of the 30-year segment. The dashed line is the 95% confidence intervals for a random collection of time series, i.e., if there were no underlying environmental cycles causing "good" and "bad" series of years. Holmes unpublished analyses.

It is tempting to attribute these 'good' versus 'bad' strings of years to an environmental driver, such as ocean conditions that one could presumably model. While this may be the case, the data by themselves do not necessarily support this since this type of cycling good and bad strings of years can happen simply by chance in a collection of stochastic population time series. Indeed this is what Figure 2 illustrates. The dotted lines in Figure 7 show the 95% confidence intervals assuming that the time series were all completely independent. This is a conservative estimate since they are not all independent and the true 95% confidence intervals are farther apart. What we can see is that the solid line falls within the conservative 95% confidence intervals suggesting this West Coast pattern of good and bad strings of years is not inconsistent with the hypothesis that it occurred by chance.

The Holmes and Fagan (2002) analysis also looked how well the diffusion approximation predicted the probability of 90% decline. This analysis searched for a difference between the mean diffusion approximation estimates of the probability of 90% decline and the observed mean probabilities within the collection of West Coast salmon time series. Figure 8 shows the estimated versus actual mean probabilities. The gray solid line (Dennis-Holmes) is the method used in the salmon PVA (McClure et al. 2003). The close correspondence between the actual and observed indicates that first the diffusion approximation approach is correctly estimating the mean probabilities and second that the parameters of this approximation were not being systematically misestimated. Note that this analysis focuses on mean estimates of probability of decline. The issue of the variability in estimates of probability of decline is addressed later in this document.
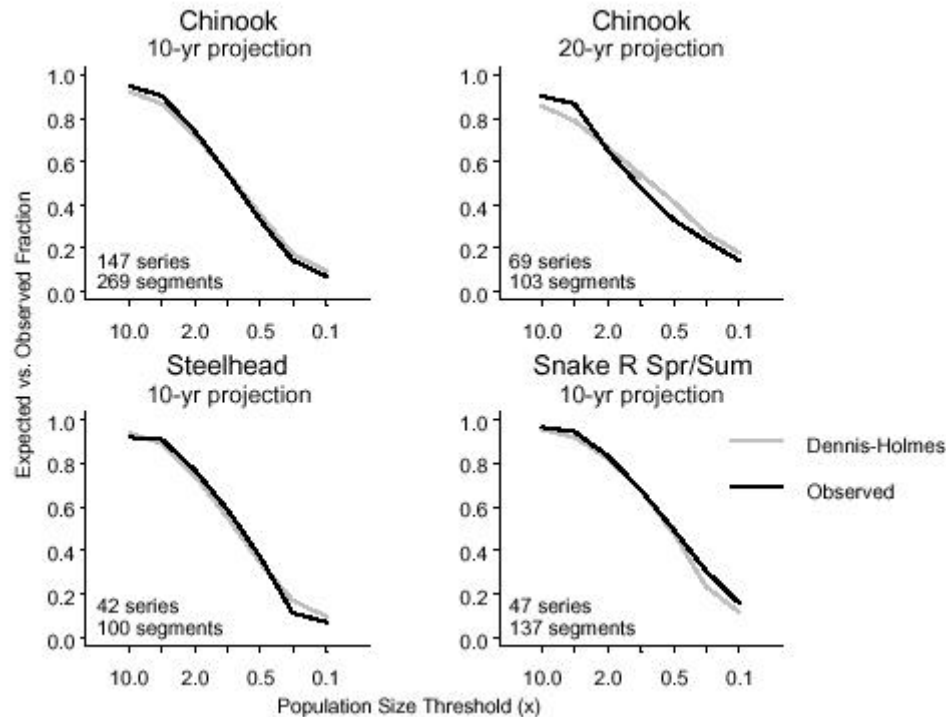
Figure 8. Probability of 90% decline versus observed probabilities with the West Coast salmon time series. From Holmes and Fagan (2002).

Why the diffusion approximation approach versus other approaches for describing trends and risks in salmon populations?

Diffusion approximation approaches for estimation of risk metrics are grounded in theoretical work on stochastic population processes (reviewed in Holmes and Fagan 2002 and Holmes 2004). These methods are one of the basic quantitative tools in population viability analysis and are featured in two current books on quantitative methods for analyzing population data (Lande et al. 2003, Morris and Doak 2003). The long-term rate of population growth is termed λ and is one of the most commonly used risk metrics within the field of conservation biology. Note that λ does not refer to a specific method of estimation, but rather simply the median or long-term trend in the population. There are a variety of methods for estimating λ. The most familiar within the conservation biology literature is to calculate λ from estimated Leslie matrix models. Diffusion approximation approaches present a way to estimate λ when only time series is available, and present a method for estimating the uncertainty in λ, which estimated Leslie matrix models do not provide.

However in the context of salmon management, traditionally other metrics of risk and population trend have been used. Some of the typical metrics that have been used or suggested are log recruits per spawner, SARs, 8-year geometric means of the natural cohort return rate, a simple regression of log natural abundance versus time, and residuals from a stock/recruit relationship. Some of these (log recruits per spawner and a regression of the log abundance versus time) have a close relationship to λ and indeed can be viewed as alternate methods for estimating λ. Many of the other methods, however, differ in a fundamental way in that they measure only a portion of the life cycle, i.e., survivorship or fecundity of only certain stages

rather than from spawner to spawner. One of the key aspects of λ is that it integrates across the entire life-cycle. It is not a measure of one stage's survivorship or fecundity alone, but rather of the integration of survivorship and fecundity over the entire life cycle, much like a spawner-to-spawner ratio does. This is important when one is trying to assess a population rather than a particular stage since high survivorship in one stage can easily be offset by low survivorship in another stage.

Below the methods that have been more common in salmon management are discussed in terms of how they relate to λ and the estimation long-term population trends.

Log recruits to the spawning grounds per spawner

Log recruits (to the spawning ground) per spawner is another way to estimate $\log(\lambda)$ since the expected value of $\ln(R/S) = \log(\lambda)$. This can be derived from theory on stochastic population processes (see review by Caswell 2001, 14.3.2) and is essentially what is shown by Eqn 14.47 in Caswell (2001) – although this probably will not be transparent on first glance. Obviously the estimates you get of $\log(\lambda)$ from Eq. 5 versus $\ln(R/S)$ are going to be different for a specific finite time series; you expect this using different methods even though the expected values (the average estimates) are identical.

If $\ln(R/S)$ can be used to estimate $\log(\lambda)$, why not use that since it is more familiar for fisheries biologists? First it is not a more accurate nor less variable estimator – a simple simulation demonstrates this. Second it requires much more data and effort to estimate – despite not providing an increase in precision in the estimation of $\log(\lambda)$. To the extent that the age-at-return data contains errors this adds additional errors to the $\ln(R/S)$ estimate. Third, if we want to compare stock status for example to prioritize recovery actions, using a consistent method across all stocks is critical. For the vast majority of stocks, the additional data to estimate R/S is not available so we can't estimate $\ln(R/S)$. Fourth, establishing the uncertainty in the estimate of $\ln(R/S)$ would be difficult. We would either have to model the error in age-at-return data, which would require some ad hoc assumptions since we have limited information on this error, or we would have to bootstrap from limited age-at-return data. Fifth, we would still have to estimate the process error and estimating this from $\ln(R/S)$ data alone is not possible if the population is affected by both process and non-process error.

8-year geometric means of the natural cohort return rate

This metric uses the 8-year geometric mean of the spawner-to-spawner ratio for the natural spawning component of the population. Like $\ln(R/S)$, this another way to estimate λ. The reasons for not using this metric are the same as those for not using $\ln(R/S)$; see above discussion.

Smolt-to-adult ratios (SARs)

SARs, along with other measures of survivorship, are clearly important for analyzing how survivorship changes within a portion of the salmon life cycle. However this metric leaves out the adult-to-smolt portion of the life cycle. For the purpose of tracking the long-term trends, the entire life cycle, spawner-to-spawner, must be included since increases in smolt-to-adult survivorship could be offset by decreases in adult-to-smolt ratios. Thus, SARs are not used for estimating long-term trends.

Note also that SARs detailed types of data, which are not available for many stocks and makes their analysis regionally limited.

Simple regression of log natural abundance versus time

λ is the regression of log spawner counts versus time for an infinite (i.e. very long) time series. One way to estimate λ is to use the regression of log spawners versus time for the available, finite, time series. This method could have been used, but simulations indicated that it gives estimates that are essentially the same as the runsum estimates (Figure 9). Even if one did use a regression, one needs to use the methods in McClure et al (2003) to get the confidence intervals on λ. The confidence intervals on the regression cannot be used since this attributes all error in the data to observation error. This is incorrect; part of the error is process error and part of it is observation error, and one needs to use a statistical framework that properly apportions the error into these two types. In addition, one still needs to obtain the estimates of environmental variability, which are critical for estimates of the probability of crossing thresholds. The regression will not provide this since again a simple analysis of the variance of the residuals attributes error to observation not process error. Holmes (2004) reviews the currently available methods for parameter estimation for population processes with process and observation error.
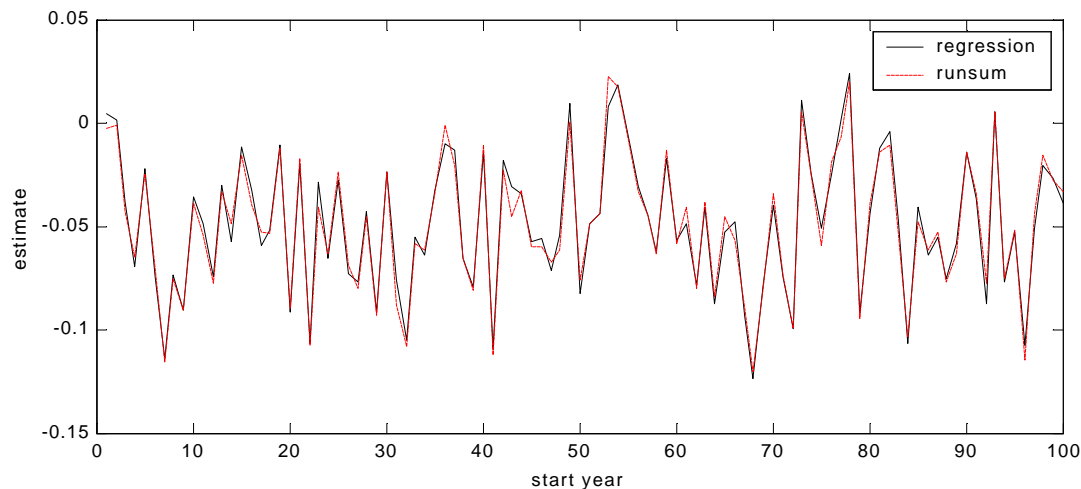

Figure 9. Example of stimated $\log(\lambda)$ from 20-year segments in a time series with $\mu = -0.05$ and $\sigma^2 = 0.02$. Segment 1 is year 1-20, segment 2 is year 2-21, segment 3 is year 3-22, etc. This shows how the estimates vary depending on the segment observed. The regression line (solid black) is from a regression of log counts versus time; the "runsum" method (red dashed) is that used in McClure et al. (2003) and for the Biological Opinion calculations.

Residuals from a stock/recruit relationship

Residuals for a stock/recruit relationship give information on how conditions in one year or cohort deviate from some longer trend. This can be useful for trying to determine if underlying changes for the long term trend has occurred, but is not useful for estimating the long term trend itself. Potentially these residuals could be used to estimate the environmental variability, although this is certainly not straight-forward. The variability in the residuals will be due not only to environmental variability but also variability due to density-dependence and the proclivity of salmon for "boom-bust" cycles. These latter types of variability are important for the short-term variability in population trajectories, but tend to dampen out with time and are less important for the long-term variability in population trajectories. Holmes (2004) gives an example of this using age-structured salmon models with density-dependence. Note also that

E. E. Holmes, National Marine Fisheries Service

residuals for a stock/recruit relationship also require age-structure data, which makes their analysis limited to stocks with that kind of detailed data.

Other methods

The methods used in McClure et al (2003) require very simple data, spawner time series, however, there are populations with much better and more detailed data, especially age-at-return and age-specific survivorship data. Incorporation of this data into the estimation of μ can increase the precision of the μ estimate, and consequently the λ estimate. Hinrichsen (2002) discusses estimation of λ using age-at-return information and shows how using this information increases precision although there is no change in bias relative to the λ estimate in McClure et al. (2003). The downside is that the methods in Hinrichsen (2002) are sensitive to high levels of observation error, for example, standard deviation of observation error greater than 0.7, which is certainly seen in redd count data (see discussion in Holmes 2004). More analytical work needs to be done to get around this sensitivity to observation error, but certainly this research suggests that more precise λ estimates can be obtained for those stocks with more extensive data. This is an area that is very promising, however, for regional analyses where we need to compare risks among stocks, some of which are data poor, we will have to continue to have and rely on methods that use only spawner time series for the sake of consistency.

Lindley (2003) presents state-space estimation for noisy time series and offers this as alternative to the estimation methods used in McClure et al. State-space estimation enables maximum-likelihood estimation of μ and σ^2 from noisy data (such as we have for salmon data). It has a strong statistical foundation. I have also been researching state-space estimation and tested Lindley's algorithm in Holmes (2004) and found that it gives much worse estimates of σ^2 than $\hat{s}^2_{slp}$ given the particular characteristics and constraints we face with salmon data. The m estimates were similar to $\hat{m}_{run}$, however. I have also investigated a slightly different state-space algorithm for estimation and found similar results. State-space estimation is extremely promising, but a significant amount of research is still need to come up with algorithms that perform more robustly than the current methods in McClure et al.

Summary

While these commonly used metrics are useful for other questions, such as looking for survivorship changes in a particular habitat or life stage or understanding the contributions of particular age classes to recruitment to the population, they are limited in terms of estimating long-term trends, either because they look at just a segment of the population, lead to λ estimates that are more variable than the λ estimates used in the Biological Opinion, or require data that is not available across all populations. Furthermore these other methods do not lead us to an estimation of underlying variability in the population process (process error), which is essential for estimation of the probability of crossing critical population thresholds and for calculating the uncertainty in our risk estimates. The methods used for estimating λ and extinction metrics as described in McClure et al. (2003) have been extensively studied and validated with West Coast salmon time series (Holmes & Fagan 2002) and also salmon-specific simulations which include density dependence (Holmes 2004).

E. E. Holmes, National Marine Fisheries Service

The time frame of ones data and λ estimates

Typically choices must be made about the data, specifically the years, to use to estimate λ. The point estimate of λ will depend on the time frame used, however keep in mind that in general the point estimate of λ should never be used alone since by itself the point estimate does not give an indication of the uncertainty in this estimate. One way present the uncertainty is to use confidence intervals, but confidence intervals are often misleading since they give the erroneous impression that the true value is equally likely within a large interval. Likelihood profiles or posterior probability distributions of λ are much more useful and give a rapid feel for the uncertainty in the estimate of λ. If one uses a posterior probability distribution, it becomes clear that the estimate of λ is not so sensitive to the time frame of the data or the addition of one extra year of data as would appear when only point estimates are presented.

This being said, selection of a reasonable time frame is very important. The following considerations should generally be kept in mind when selecting the time frame to use: a) more data is better, b) the time frame should be representative of historical trends, i.e. not be dominated by 'good' or 'bad' conditions and not dominated by an isolated perturbation and c) for the sake of uniformity and comparison, the time frame should be consistent across stocks. In McClure et al. (2003) the effect of using different time frames for estimation, specifically 1980-2000 versus 1960-2000, on risk metrics for the Columbia River ESUs is shown. The differences were not statistically significant nor in any consistent direction, i.e. for some stocks the 1980-2000 time period gave slightly more severe risk estimates and for others it gave less severe estimates.

From a management standpoint, λ estimates that vary widely depending on the exact starting year of the time series are problematic, and research showing that the estimates are statistically optimal while satisfying does not lessen this practical problem. There are a couple of strategies that I have proposed to deal with this:

1) Use robust estimators of the mean for the μ estimates. Currently in Eqn. 5, a straight mean is used, however a straight mean is highly sensitive to outliers. My preliminary tudies of the effect of different start years on λ estimates using Snake River spring/summer chinook time series indicated that a robust estimator of the mean eliminated much of the problem of λ estimates that vary widely depending on the start year. There are a variety of robust mean estimators; a trimmed mean is the simplest.

2) I examined the 1970s to present data throughout the Columbia River and found that the 1980-present data was affected by an especially unusual series of years between 1978-1982 or so. The estimates using the 1980-present time frame appeared to be more different that one would expect compared to estimates using any other time frame. My initial analysis suggested that 1976-present would generally be a better time frame to use, although this does suffer from dam effects in the early years for some stocks. The 1984-present data could also be used to avoid the 1978-82 period, however, a strong argument can be made that this overly emphasizes a period characterized by bad ocean conditions.

II. ACCOUNTING FOR UNCERTAINTY IN RISK ESTIMATES

A certain amount of variability in estimated parameters and risk metrics is an unavoidable aspect of the analysis of stochastic population processes, simply due to the nature of these processes. One of the strengths of diffusion approximation methods is that the statistical distributions of the estimated parameters are known. As a result, the uncertainty in the estimated risks can be calculated. This is often not the case for other PVA approaches. Even though the

uncertainty in diffusion approximation risk metrics can be calculated, this uncertainty is definitely high. In this situation, examining either the likelihood functions or the posterior probability distributions for the risk metrics, rather than simply the point estimates and confidence intervals, will help to clarify the level of data support for different true risk levels. Statistical decision theory (e.g. Berger 1985 is one of many texts on decision theory) provides a framework for integrating estimates of the data support for different risk levels with the consequences of true risk levels. Wade (2000) and Dorazio and Johnson (2003) provide recent discussions of this Bayesian decision framework in conservation biology and resource management contexts.

The idea in a nutshell is to estimate the probability that the risk metric, for example λ, is within particular ranges that are important from a management perspective. For example:

Probability	< 0.9	0.15
the true λ is	0.9 – 0.95	0.3
in these	0.95 – 1.0	0.5
ranges	> 1.0	0.05

Table 1. Estimated probabilities that λ is within different ranges.

These probabilities are estimated using the posterior probability distribution that is estimated from the data. Figure 10 gives an example of the posterior probability distribution for λ estimated from a 38-year times series of spring chinook in the Upper Columbia River basin (data from T. Cooney, NMFS). The probability that λ is within the range *a* to *b* is calculated by integrating the posterior probability distribution between *a* and *b*.
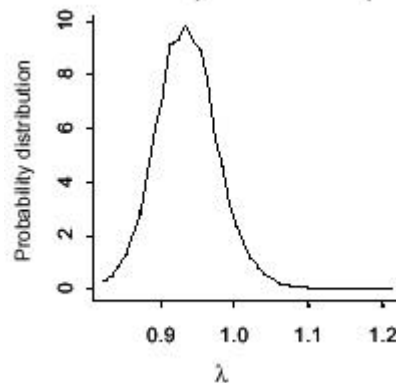

Figure 10. Estimated posterior probability distribution for λ for Upper Columbia spring chinook. From Holmes (2004).

There are a variety of ways these probabilities could be used. They might be used alone and qualitative thresholds set, such as if the probability that λ is less than 0.95 is greater than some threshold, then an such and such actions can (or cannot) occur. Note that the probability of a low λ can be high due to certainty that λ is low or due to uncertainty about λ. Thus, such a strategy leads to caution in the face of high uncertainty. A more quantitative, decision-theoretic, approach can be taken if the probability that actions will be 'sufficient' (however that might be defined) can be calculated given different true λs. For example,

E. E. Holmes, National Marine Fisheries Service

Probability of action being 'sufficient'			
	Action A	Action B	Action C
< 0.9	0.3	0.1	0
0.9 –0.95	0.5	0.3	0.2
0.95-1.0	0.8	0.6	0.5
> 1.0	1.0	1.0	0.8

λ range is a label applied to the left of the four λ range rows.

Table 2. Estimated probabilities of action sufficiency given different true λ ranges.

These probabilities are multiplied by the probability of λ being within those ranges and then summed over all ranges to give the total probability that actions are 'sufficient'. This probability incorporates the uncertainty in the estimated λ:

Probability of action being 'sufficient'		
Action A	Action B	Action C
0.64	0.45	0.35

Table 3. Probabilities in Table 2 multiplied by those in Table 1 and summed over all λ ranges.

An example where the probabilities in Table 2 would be relatively easy to calculate is different harvest levels. Instead of giving a simple 'yes/no' answer as would be the case if using point estimates of λ, this approach quantifies the uncertainty in our estimate of λ and emphasizes that there is not a simple "100% or 0%" probability of an action being effective.

Probabilities of crossing thresholds are notoriously uncertain and variable, and analyzing the uncertainty connected with a proposed probability metric (e.g. 'probability of extinction') is especially critical when using these metrics. Figure 11 shows the estimated probability density distributions for the probability of 90% decline within 25, 50 or 100 years given a 20-year time series with an estimated λ of 1 or 0.93. The distributions when the estimated λ is 1 are fairly flat or U-shaped. This indicates that there is not much information about what the probability of 90% decline is. The estimation of the probability of 90% decline can be improved by using an informative prior on the process error. Twenty years of data is not sufficient for accurate process error estimates. If one argues that the variability driving long-term dynamics is similar across chinook throughout the basin, then one might use as an informative prior the distribution of process error estimates for a large number of stocks throughout the basin. Figure 12 shows how the estimation of the probability of 90% decline improves using an informative prior. Now it appears that estimation of the risk of 90% decline in 50 or 100 years is fairly informative for the stock with a low λ. For the stock with a λ equal to 1, 50 and 100-year probabilities are uncertain, but 25-year probability are much better.
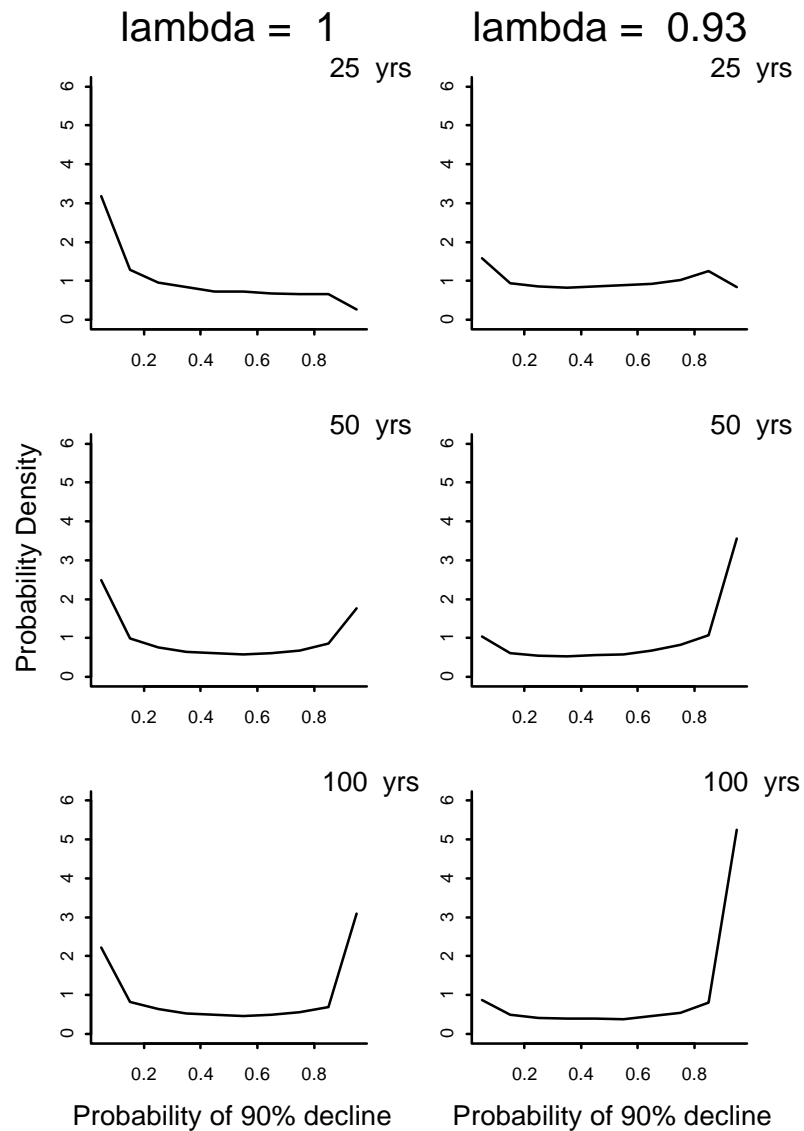
Figure 11. Estimated posterior probability distribution of the probability of 90% decline in 25, 50 and 100 years given a 20 year time series with estimated μ of 0 or –0.072 and an estimated process error of 0.08 and estimated non-process error of 0.71. A uniform prior on the process error was assumed.
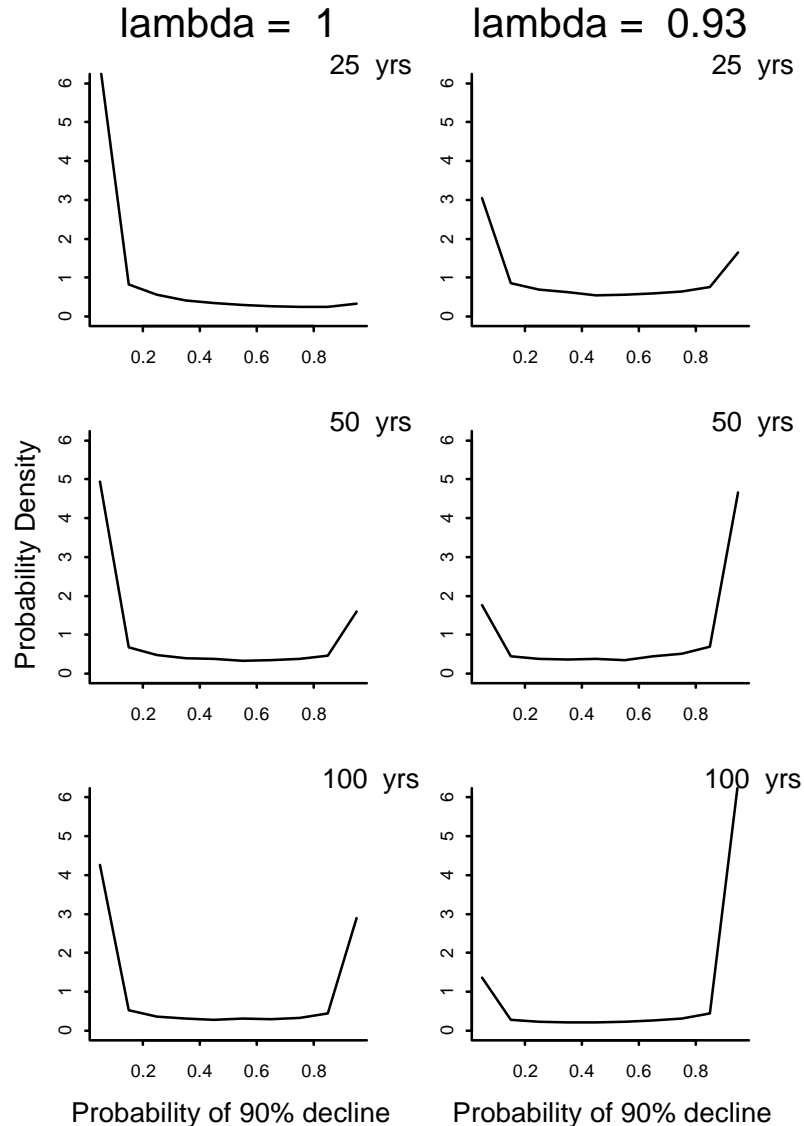
lambda = 1   lambda = 0.93

Figure 12. Estimated posterior probability distribution of the probability of 90% decline in 25, 50 and 100 years given a 20 year time series with estimated μ of 0 or –0.072 and an estimated process error of 0.08 and estimated non-process error of 0.71. A highly informative prior on the process error was assumed.

These last two figures focus on the probability of 90% decline in 25, 50 or 100 years. There are other ways to look at the probability of 90% decline that can be more informative. For example, here is an analysis of the probability of an eventual 90% decline for based on the Upper Columbia spring chinook time series (from Holmes 2004). The estimated probability of eventual 90% decline is almost 1.0, that is it is almost certain to occur (Figure 13, top), however there a great deal of uncertainty as to when this will occur (Figure 13, bottom) except that it is highly likely within 100 years. Figure 13, bottom panel, shows the expected probability of 90% decline within a given time frame. From the figure, on average there is a 70% probability of a 90% decline within 50 years for this population and an average 80% probability that the 90% decline occurs within 100 years.
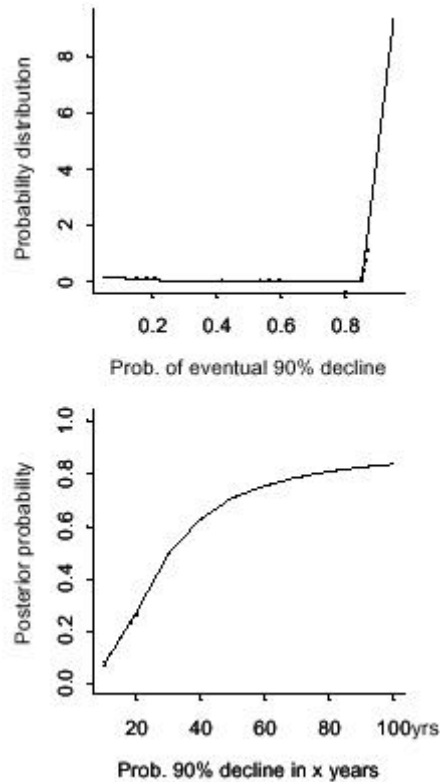
Figure 13. Estimated posterior probability distribution of eventual 90% decline in the Upper Columbia spring/summer chinook (top panel) and the probability that this decline has occurred within different time frames (bottom panel). From Holmes (2004).

The time to 90% decline can then be put into Table form similar to Table 1 for λ ranges (Table 4) which can then be combined with a Table similar to Table 2 for the probability that an action will be sufficient if 90% decline does not occur until after x years.

Probability	> 25 yrs	0.7
90% decline	> 50 yrs	0.38
only after x	> 75 yrs	0.2
years	> 100 yrs	0.18

Table 4. Probability that 90% decline does not occur until after x years. From Figure 13, bottom panel.

Quasi-extinction threshold versus absolute extinction

Throughout this paper, the probability of 90% decline is discussed. The probability of 90% can be estimated for stocks for which we only have index data and not total spawner information. Thus it can be more widely applied. However, decline to specific critical population sizes are also of great importance in PVA analyses. Although estimating extinction to 1 individual is a popular risk metric, and unfortunately sometimes mandated, caution is required when using the diffusion approximation is to estimate extinction to very low numbers since factors that drive dynamics at very low population sizes (such as demographic stochasticity) and the catastrophes often associated with ultimate extinction will likely be poorly

E. E. Holmes, National Marine Fisheries Service

represented in a time series of a relatively larger population declining to low numbers. There have been a wide variety of papers published on this in the conservation biology literature. The general recommendation is to estimate the probability of decline to some critical population size (quite a bit greater than 1); this is termed a 'quasi-extinction' threshold. Fagan et al. (unpublished manuscript) studied a collection of actual time series of species that went extinct and compared diffusion approximations for quasi-extinction thresholds versus extinction to 1 individual. This analysis found that quasi-extinction estimates (to a size much greater than 1) using diffusion approximations fit the observed data, but that extinction estimates (to 1 individual) were very poor and underestimate the true risk.

With this in mind, one might ask why was the probability of decline to 1 individual estimated in the McClure et al. analysis. The reasons for this were four-fold. 1) The analysis was focused on estimating risks if current conditions continue as they appeared in the time series data. It was recognized that this would tend to underestimate risks if factors such as density-depensation occurred as the population got small, however baseline estimates of risks under current conditions are required in order to make meaningful statements about risks under hypothetical future scenarios, such as lower population growth rates as the population gets small. 2) The probability of 90% decline does not incorporate the actual population size. The implications of a 90% decline of a population of 10 individuals is quite different than a 90% decline for a population of 100,000 individuals. The probability of decline to 1 individual provided a risk metric that incorporated both the overall rate of decline of the population and the population size. Thus we could then compare ESUs in terms of a risk metric that integrates these two factors –regardless of whether this is an underestimate of the true probability of extinction. 3) Any other extinction threshold we could have specified would have been arbitrary – given the information we had on critical population size. Decline to 1 individual is meaningful for all populations. 4) The Fagan et al. analysis had one notable exception, i.e. one population where the probability of extinction would be properly estimated. That was the one population time series that followed an actual salmonid extinction (sockeye); the rest of the time series followed bird and reptile extinctions.

REFERENCES

Achord, S., P. S. Levin, and R. W. Zabel. 2003. Density-dependent mortality in Pacific salmon: the ghost of impacts past? Ecology Letters 6: 335-342.

Berger, J. O. 1985. Statistical decision theory and Bayesian analysis. Springer Verlag, NY, NY, USA.

Caswell, H. 2001. Matrix Population Models. Sinauer, Sunderland, Mass.

Dennis, B., P. L. Munholland, and J. M. Scott. 1991. Estimation of growth and extinction parameters for endangered species. Ecological Monographs 61:115-143.

Dorazio, R. M., and F. A. Johnson. 2003. Bayesian inference and decision theory —a framework for decision making in natural resource management. Ecological Applications 13:556–563.

Fagan, W. F., J. Rango, A. Folarin, J. Sorensen, J. Lippe, and N. E. McIntyre. Cross-validation of quasi-extinction risks from real time series: an examination of diffusion approximation methods. Manuscript.

Gerber, L., D. DeMaster and P. Kareiva. 1999. Grey whales and the value of monitoring data in implementing the U.S. Endangered Species Act. Conservation Biology 13:1215-1219.

Hinrichsen, R. A. 2002. The accuracy of alternative stochastic growth rate estimates for salmon populations. Canadian Journal Fisheries and Aquatic Sciences 59:1014–1023.

Holmes, E. E. 2001. Estimating risks in declining populations with poor data. Proceedings of the National Academy of Science USA 98:5072-77.

Holmes, E. E. and W. F. Fagan. 2002. Validating population viability analysis for corrupted data sets. Ecology 83: 2379-2386.

Lande, R. and S. H. Orzack. 1988. Extinction dynamics of age-structured populations in a fluctuating environment. Proceedings of the National Academy of Science USA 85:7418-7421.

Lande, R., S. Engen, and B. Saether. 2003. Stochastic population models in ecology and conservation: an introduction. Oxford University Press, Oxford, UK.

Lindley, S. T. 2003. Estimation of population growth and extinction parameters from noisy data. Ecological Applications 13: 806-813.

McClure, M. M., E. E. Holmes, B. L. Sanderson, and C. E. Jordan. 2003. A large-scale, multi-species risk assessment: anadromous salmonids in the Columbia River Basin. Ecological Applications 13:964-989.

Morris, W. F. and D. F. Doak. 2002. Quantitative conservation biology: theory and practice of population viability analysis. Sinauer Press, Sunderland, MA.

Morris, W., D. Doak, M. Groom, P. Kareiva, J. Fieberg, L. Gerber, P. Murphy, and D. Thomson. 1999. A practical handbook for population viability analysis. The Nature Conservancy.

Nicholls, A. O., P. C. Viljoen, M. H. Knight and A. S. Van Jaarsveld. 1996. Evaluationg population persistence of censused and unmanaged herbivore populations from the Kruger National Park, South Africa. Biological Conservation 76:57-67.

Tuljapurkar, S. D. 1989. An uncertain life: demography in random environments. Theoretical Population Biology 35:227-294.

Tuljapurkar, S. D. and S. H. Orzack. 1980. Population dynamics in variable environments. I. Long-run growth rates and extinction. Theoretical Population Biology 18:314-342.

Wade, P. R. 2000. Bayesian methods in conservation biology. Conservation Biology 14: 1308-1316.